

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: UPDATING DATA IN A MULTI-SYSTEM NETWORK
THAT UTILIZES ASYNCHRONOUS MESSAGE
TRANSFER

APPLICANT: FRANK MICHAEL KRAFT AND GUENTER PECHT-
SEIBERT

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EV 342626655 US

December 2, 2003
Date of Deposit

Updating Data In A Multi-System Network that Utilizes Asynchronous Message Transfer

TECHNICAL FIELD

This invention relates to asynchronous data transfer between different computing systems, and more particularly to techniques for updating data that is maintained and modified by two or more computing systems in a multiple-system environment where asynchronous messaging is employed.

BACKGROUND

A multiple-system network may include two or more systems that share a common data set. The multiple systems may share the data set by maintaining a local version of the data set in a local database. In some networks, each of the multiple systems may have the ability to modify or change the data set. These changes may be made to the system's local version of the data set and then communicated to the other systems in the network via asynchronous messaging. The messages allow the other systems to update their local data sets with the modifications that occurred in other systems so that the data sets of the various systems are consistent. Because asynchronous messaging may be used to update the various data sets in the network, it is possible that messages between the systems could be delayed or even lost. In the event of a lost message, a state of equilibrium between the data sets of the system may not be reached.

SUMMARY

The invention provides techniques for updating a locally stored version of a data set, wherein multiple application systems each maintain a separate stored version of the data set and are each capable of modifying an attribute of the data set. The attribute comprises multiple attribute components that are each associated with a different one of the multiple application systems, and each attribute component has a data value and a change value.

In an aspect, the invention provides a method for performing such an update. In the method, a message containing a data set as locally stored in a first application system is received at a second application system of the multiple application systems. For each attribute component of the data set, the attribute component's change value as stored locally in the second application system is compared with the attribute component's change value as contained in the received data set. If the comparison indicates that the version of the attribute component's data value as stored locally in the second application system is less recent than the version of the attribute component's data value contained in the received data set, the attribute component's data value stored locally in the second application system is replaced with the attribute component's data value contained in the received data set.

In another aspect, the invention provides a method for updating a locally stored version of a data set. In the method, an attribute of a data set stored locally in a first application system is modified. A message is generated containing the modified data set as locally stored in the first application system, the data set comprising multiple attribute components that are each associated with a different one of the multiple application systems. The message is sent to a second application system that maintains a version of the data set.

In embodiments, the methods may have one or more of the following features. The change value of the attribute component may be a timestamp that indicates the time of the modification to the attribute component's data value. The change value of the attribute component may also be a version number that is incremented after each modification to the attribute component's data value. The message may also have a total value that represents a sum of the attribute component data values in the data set as locally stored in the first application system. In some implementations, the message received by the second application system from

the first application system may also include the attribute components associated with the second application system.

In some embodiments, the first application system may send messages to the second application system after each modification of the first application system's data set. The first application system may send a message to the second application system containing the first application system's data set after each modification of an attribute component of the first application system's data set. Asynchronous messages may be used to send the messages to the second application system.

In other aspects, the invention provides computer program products that perform the above-described methods. In particular, the computer program products comprise executable instructions embodied on either a computer readable medium or a propagated signal.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

15

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a multiple-system network in which a data set is maintained and modified by the multiple systems.

FIG. 2 is a diagram showing an example protocol for a message that may be transferred between the systems of FIG. 1.

20

FIG. 3 is a block diagram of a computer system that may be included in the systems of FIG. 1.

FIG. 4 is a flowchart of a method of updating a data set of a system shown in FIG. 1.

FIG. 5 is a table that illustrates the method of updating the data sets of the multiple systems shown in FIG. 4.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

5 A multi-system network 10, shown in FIG. 1, includes three networked computing systems, which in this example are a first system 20, a second system 30, and a third system 40. The systems 20, 30, and 40 each maintain a data set that is shared by the systems in the network 10. Each system's data set includes multiple data values and change values that correspond to 10 each data value. The change values track the modifications or changes that are made to the data values. For each modification made by one of the systems 20, 30, and 40 to a data value in the data set, the modifying system sends an asynchronous message containing the system's entire data set to the other systems in the network 10. When a system receives the message, the system compares the change values in the received message to the corresponding change values in the 15 system's local data set. If the change values indicate that a particular data value in the message is more recent than the corresponding data value in the system's data set, the system updates its local data set with the more recent data value in the message. By including in the update message all of the data values and change values in the modifying system's data set, any inconsistencies that may exist among the data sets of the multiple systems in the network 10 due to delayed or lost messages may be repaired.

20 The first system 20 includes a database 22, a user interface 24, and a message transport layer 26. The database 22 contains the system's local version of the data set that is shared by the various systems 20, 30, and 40 in the network 10. In the FIG. 1 example, the data set relates to products or inventory information and may be used to track the quantities of items that are

purchased, sold, returned, etc. In other examples, the data set may relate to banking information, accounting information, or other types of information that may be maintained using the methods described herein.

The data set in this example includes a data value for its own system 20 and a data value
5 for each of the other systems 30 and 40 in the network 10. In other examples, the data set may include multiple data values for the systems, or alternatively, the data values for certain systems may be omitted. The data set also includes a corresponding change value for each of the data values. The change value may be, for example, a version number that is incremented after each change or modification of the data value. Alternatively, the change value may be a timestamp
10 that reflects the actual time of the modification. Yet in other implementations, a different non-cyclical identifier may be used to track the changes to the data values.

The data values in the data set may be changed or modified through the user interface 24. These changes may be made by a system administrator accessing the first system 20 or,
alternatively, by other systems in the network that are authorized by making these modifications.

15 When a data value is modified in the database 22, the first system 20 sends a message to the other systems 30 and 40 in the network 10 to update their respective versions of the data set. In this example, the messages are sent to the other systems asynchronously; however, the messages may also be transmitted between the systems using other suitable methods. These messages are sent to the other systems in the network through message transport layer 26. The message
20 transport layer 26 also receives messages from the other systems. In this example, a network 12 connects the first system 20, second system 30, and third system 40. The network 12 may be a LAN, the Internet, or another suitable network.

In the FIG. 1 example, the second system 30 and third system 40 are similar to the first system 20. The systems 30 and 40 include databases 32 and 42 that store the each system's

version of the data set, user interfaces 34 and 44 through which each system may modify its data set, and message transport layers 36 and 46 to send and receive messages.

FIG. 2 is a diagram of an example message format that may be used for messages transferred between the systems 20, 30, and 40 of FIG. 1. The message 50 includes a sending system identifier 52, which identifies the system from which the message is being sent. A delta value is also included in the message 50. The delta value indicates the value of the change that was made to the sending system's data set that prompted the generation of the message 50. For example, the data value would be +5 if the sending system modifies its data value from five to ten. In other implementations, the delta value 52 may be omitted.

The message 50 also includes a first system data value 56, a second system data value 58, and a third system data value 60. The data values 56, 58, and 60 reflect the most recent data values for the respective systems in the sending system's data set. The data value may be a number, such as in the inventory example, or in other examples, may be some other data type. The message also includes first system change value 62, second system change value 64, and third system change value 66. The change values 62, 64, and 66 correspond to the data values 56, 58, and 60. The change values 62, 64, and 66 allow the system receiving the message 50 to determine whether the receiving system's data set includes the most recent version of the data values, as is discussed in greater detail below. The message also includes a total value 68. In the FIG. 1 example, the total value 68 represents the sum of the system data values 56, 58, and 60. In other examples, the total value 68 may represent other information or may be omitted.

FIG. 3 is a block diagram of a computer system 70 that may be included in the systems 20, 30, and 40 of FIG. 1. The computer system 70 includes a program memory 72 containing a message program 74 and an updating program 76. The message program 74 contains instructions that when loaded into RAM 80 and executed by processor 78 generate a message for

transmission to another system in the network (for example, the message shown in FIG. 2). The processor 78 may obtain the information in the FIG. 2 message from the system's data set 82. The message may then be stored in RAM 80 until it is output by an input/output module 84 to a message transport layer (shown in FIG. 1). Messages may also be received through the 5 input/output module 84 and stored in RAM 80 for processing. The updating program 76 contains instructions that when loaded into RAM 80 and executed by processor 78 processes the received messages stored in RAM 80 and, if necessary, updates the data values stored in data set 82.

FIG. 4 is a flowchart of a method 100 of updating a data set of a system. The FIG. 4 10 example illustrates the method 100 using the example of FIG. 1. The method 100 begins at step 110 with the modification of the first system's data set, as described previously. Next, the first system 20 sends an asynchronous message to the other systems 30 and 40 containing the first system's modified data set at step 120. The message may contain a data value for each of the 15 systems 20, 30, and 40 and corresponding change values, such as the message 50 shown in FIG. 2. The message may also include a sending system identifier, a delta value, and a total value; however, these message components are optional and may be omitted.

The second system 30 receives the update message sent by the first system 20 at step 130. After the message is received, the second system 30 compares the change values in the message with the corresponding change values in the second system's data set at step 140. If the second 20 system 30 determines at step 150 that each data value in second system's data is as recent or more recent than the corresponding data values in the message, the second system ignores the contents of the message at step 160. If, on the other hand, the change values indicate that a data value in the message is more recent than a corresponding data value in second system's data set, the data value in the second system's data set is replaced with the more recent data value in the

message at step 170. For each data value in the second system's data set that is replaced with a data value in the message, at step 180 the second system replaces the corresponding change value in the second system's data set with the change value in the message.

In other examples, the sending system identifier (shown in FIG. 2) may be included in the update message. For example, in implementations where the data value and change value of the system receiving the update message are not included in the update message, the message may include the sending system identifier to indicate that the receiving system's data value and change value were intentionally left out of the message.

In other implementations, the total value (shown in FIG. 2) may be included in the update message. The total value may be used to verify a state of equilibrium between the data sets of the sending and receiving systems. For example, if the change values in the message and the change values in the receiving system's data set are equal after the receiving system's data set is updated, the total value in the message and the data set may be compared. If the total values are equal, then the data sets of the sending and receiving systems are in a state of equilibrium. If the total values are not equal, then another process may be run by the systems to repair the inconsistencies. Yet in other examples, the total value may be used to detect errors in the data values transmitted in the message.

FIG. 5 is a table that illustrates the FIG. 4 method of updating the data set of a system in a network. The table 200 includes three columns 202, 204, and 206 that correspond to the first system 20, second system 30, and third system 40, respectively. Each row of the table 200 represents a time periods from T0 to T16. The contents of the data sets shown in the columns 202, 204, and 206 are shown in a format +A (B, C, D)[E, F, G] H. This format is consistent with the example message protocol shown in FIG. 2. For example, the +A represents the delta value. The data value vector (B, C, D) represents the first system data value (B), the second system data

value (C), and the third system data value (D). The change value vector [E, F, G] represents the first system change value (E), the second system change value (F), and the third system change value (G). The H represents the total value, or the sum of the components of the data value vector. Using the FIG. 1 inventory example, the data vector includes numbers that represent the 5 cumulative total of the inventory modifications made by each of the systems in the network. The change value in this example is a version number that is incremented after each modification of the corresponding data value. Thus, in this example, a higher change value number indicates that the corresponding data value is more recent. The total value represents the total amount of inventory in the combined systems.

10 At T0, the contents of the data sets for the systems are in a state of equilibrium at $(0, 0, 0)[0, 0, 0]$. At T1, the data set of the first system in column 202 is modified by the addition three units of inventory. The contents of the first system's data set is changed to $(3, 0, 0)[1, 0, 0]$ 3, which indicates that the first system data value is three, the first system's change value is one, and the total value is three. After the modification of the first system's data set, the first system 15 sends an asynchronous message to the other systems containing the modified data set, as discussed previously.

At T2, the first system's data set is modified again with the addition of four more units of inventory. The contents of the first system's data set is changed to $(7, 0, 0)[2, 0, 0]$ 7, which indicates that the first system's data value is seven, the first system's change value is two, and the 20 total value is seven. After the modification, the first system sends another asynchronous message to the other systems in the network.

At T3, the second system receives the T2 message $(7, 0, 0)[2, 0, 0]$ containing the first system's modified data set. In the FIG. 5 example, the update messages only include the data values and the change values. In other examples, the sender identifier, the delta value, and the

total value may also be included, as discussed previously. In accordance with the FIG. 4 method, the second system compares the change values of the message to the corresponding change values of the second system's data set. The comparison reveals that the first system change value in the message (2) is more recent than the corresponding change value of the data set (0).

5 Therefore, the second system replaces the first system data value in the data set with the first system data value in the message. The corresponding first system change value is also replaced. After the update, the contents of the second system data set are (7, 0, 0)[2, 0, 0] 7.

At T5, the data set of the third system is modified with the addition of two units of inventory. With this modification, the third system's data set is changed to (0, 0, 2)[0, 0, 1] 2. At 10 T6, the data set of the second system is modified with the subtraction of four units of inventory. The data set of the second system (7, 0, 0)[2, 0, 0] 7 is updated to (7, -4, 0)[2, 1, 0] 3. After the modifications, both systems send their modified data set to the other systems in the network.

At T7, the third system receives the update message (7, -4, 0)[2, 1, 0] generated by the second system after the modification at T6. The contents of the message are compared to the 15 contents of the third system's data set (0, 0, 2)[0, 0, 1] 2. The change values for the first and second systems indicate that the first system data value (7) and the second system data value (-4) are more recent than the corresponding data values of the third system data set. At T8, the data values and corresponding change values in the third system data set are replaced in accordance with the FIG. 4 method to (7, -4, 2)[2, 1, 1] 5. In this example, the third system did not receive 20 either of the update messages generated by the first system at T1 or T2. Because the update from the second system included the data values and change values for all systems in the network, the third system was updated with the first system modifications previously received by the second system.

At T9, the first system receives the update message $(0, 0, 2)[0, 0, 1]$ generated by the third system after the modification at T5. At T10, the contents of the first system's data set are compared to the contents of the T9 message. The comparison indicates that the data value for the third system is more recent than the data value in the first system's data set. Accordingly, the 5 first system's data set $(7, 0, 0)[2, 0, 0] 7$ is updated to $(7, 0, 2)[2, 0, 1] 9$.

At T11, the third system receives the message $(3, 0, 0)[1, 0, 0]$ generated by the first system after the modification at T1. The first system change value in the third system's data set indicates that the data set's first system data value is more recent than the corresponding data value in the message. The third system's data set is more recent because the modifications of the 10 first system at T1 and T2 were received through an update message from the second system at T7. Consequently, the values in the third system's data set are more recent and the T1 update message is ignored at T12.

At T13, the first system receives the message $(7, -4, 0)[2, 1, 0]$ generated by the second system after the modification at T6. The second system change value in the message is more 15 recent than the second system change value in the first system's data set. Accordingly, at T14 the first system data set $(7, 0, 2)[2, 0, 1] 9$ is updated with the second system data value in the message and the corresponding change value to $(7, -4, 2)[2, 1, 1] 5$.

At T15, the second system receives the message $(0, 0, 2)[0, 0, 1]$ generated by third system after the modification at T5. The message's third system change value is more recent 20 than the third system change value of the second system's data set. Thus, the second system data set $(7, -4, 0)[2, 1, 0] 3$ is updated with the third system data value in the message and corresponding change value to $(7, -4, 2)[2, 1, 1] 5$. After this modification, the data sets of the three systems are in equilibrium as shown in the last row of the table 200 labeled "status."

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, the method of updating the data sets of the systems is applicable to networks with more than three systems. The method may also be used in networks that may develop inconsistencies between the various system data sets but do not employ asynchronous messaging. Accordingly, other embodiments are within the scope of the following claims.

5